Hedging Against the Silver Tsunami: A Health Capital Preservation Triage Framework to Mitigate Long-Term Care Liability and Insurance Risks

Shiming Wu

January 2026

Abstract

As populations age, the primary economic threat to healthcare sustainability is the accelerated depreciation of health capital. While traditional triage focuses on short-term liability shielding, musculoskeletal (MSK) triage in elderly populations requires a long-term capital preservation approach. This research proposes three novel frameworks: The Functional Autonomy Hedge (FAH), which utilizes gradient-based symptom volatility; Expected Liability Calibration (ELC), which optimizes for insurance solvency by pricing the marginal cost of diagnostic delay; and the Frailty-Decay Shield (FDS), which incorporates a biological aging parameter. By identifying latent invisible risks near clinical cliffs, I employ actuarial policies to prioritize patients with high-risk profiles. Simulation results ($N = 5,000$) demonstrate that while FAH offers a modest 5.6% liability reduction over standard AI, the optimized ELC framework achieves a 50.2% reduction. The FDS framework, by pricing frailty, provides the superior result with a 52.4% reduction, effectively serving as a mathematically rigorous hedge against the systemic risks of the Silver Tsunami.

1 Introduction

The global healthcare landscape is currently confronting a demographic shift colloquially termed the "Silver Tsunami", a phenomenon where the proportion of the population over age 65 is projected to double by 2050. Within the Canadian context, Musculoskeletal (MSK) disorders have emerged as the leading cause of years lived with disability (YLDs), placing unprecedented strain on clinical resources and insurance solvency. Unlike acute cardiac or trauma events, MSK pathology often lacks objective biomarkers in its early stages. Instead, it manifests as subjective pain narratives that are frequently misunderstood or dismissed by front-line digital triage tools like HealthLink 811. This creates the phenomenon of "Invisible Pain": a diagnostic gap where patients with precarious neurological or metabolic conditions are routed to standard physiotherapy waitlists, where their functional status deteriorates during the delay.

The central research question of this paper is: Can we reframe triage from a static clinical classification task into a dynamic health capital preservation problem to mitigate the economic and clinical costs of diagnostic delay? While current triage algorithms optimize for point-estimate accuracy, they fail to account for the marginal rate of functional decline. In an aging society, a patient is not merely a data point in a queue but a depreciating health asset. This paper investigates whether quantifying symptomatic sensitivity—how close a patient sits to a clinical cliff—can allow insurers to hedge against the catastrophic financial liabilities associated with long-term care (LTC) admission and medical malpractice.

Solving this research question is of critical importance to both healthcare providers and the insurance ecosystem. From a clinical perspective, MSK disorders are the primary Gateway to Dependency. A delayed diagnosis of spinal cord compression or metastatic bone disease leads to an irreversible loss of functional autonomy, often forcing an elderly patient out of independent living and into high-cost institutional care. From an actuarial perspective, these events represent fat-tail risks. A missed metabolic red flag does not merely represent a clinical error; it triggers a multi-million dollar insurance event. By pricing the risk of delay rather than just the likelihood of disease, we can transform triage from a passive gatekeeping mechanism into an active tool for protecting health capital and ensuring the solvency of public and private insurance risk pools.

To address this, I propose the Actuarial Patient Advocacy (APA) framework, which operates as a two-stage engine. In the first stage, I employ a Deep Multi-Layer Perceptron to map eight subjective and functional markers—such as night pain severity and walking loss—into four specialist pathways: Mechanical, Inflammatory, Neurological, and Metabolic. In the second stage, we move beyond static prediction by implementing two distinct triage ranking methodologies. The Functional Autonomy Hedge (FAH) utilizes backpropagation to calculate a Volatility Index, measuring the sensitivity of the patient's risk to marginal symptom changes. Simultaneously, the Expected Liability Calibration (ELC) methodology reframes the queue as a portfolio optimization problem, pricing each patient based on the expected insurance cost of their specific pathology (e.g., \$2M for a missed metabolic diagnosis vs. \$150k for a neurological delay).

My experimental results, derived from high-fidelity simulations of 5,000 patients, demonstrate a breakthrough in risk mitigation. My full framework, incorporating three distinct methodologies, achieves a clear hierarchy of performance. The ELC methodology identified an optimal actuarial equilibrium at a hedging intensity of $\lambda = 0.91$, resulting in a \$34.8 million (50.2%) reduction in liability compared to a standard risk-neutral AI protocol. The superior Frailty-Decay Shield (FDS) framework achieves a total liability reduction of over \$36.3 million (52.4%). These findings prove that reframing triage as a multi-stage actuarial hedge provides a mathematically rigorous pathway to securing healthcare sustainability in the era of the Silver Tsunami.

This research synthesizes three distinct domains to create a novel triage paradigm for an aging demographic: Clinical Pathways, Safe AI, and the Health Economics of Aging.

First, this paper contributes to literature of Clinical Triage & MSK Pathways. Conventional MSK triage relies heavily on static heuristics for patient categorization. Hill et al. (2011) established the efficacy of the STarT Back Tool for primary care stratification, which remains a benchmark for low-risk management. However, Foster et al. (2013) argues that while stratified models of care are a significant advancement, the practical implementation often struggles to account for the multidimensional complexity and varying clinical trajectories of individual patients. Furthermore, Newman-Toker et al. (2019) highlights that diagnostic errors in neurological and systemic presentations contribute disproportionately to severe patient harm. My APA framework addresses these limitations by utilizing a deep neural engine to map subjective narratives into four distinct pathways, effectively closing the diagnostic gap for "Invisible Pain" and latent pathologies that remain undetectable by traditional heuristic checklists.

Moreover, this paper builds on papers about Bayesian Deep Learning & Algorithmic Safety methodology. The quantification of model confidence is essential for safety-critical healthcare. While Gal and Ghahramani (2016) provided the mathematical foundation for approximating Bayesian uncertainty using MC Dropout, and Leibig et al. (2017) demonstrated its efficacy in disease detection, these methods generally treat uncertainty as a passive flag for human review. Moreover, Obermeyer et al. (2019) warned that algorithms focused strictly on cost proxies can inadvertently marginalize complex patients with under-coded needs. The Functional Autonomy Hedge (FAH) proposed in this paper solves this by transforming gradient-based volatility from a passive metric of model confusion into an active triage signal. By prioritizing patients situated near clinical cliffs, I ensure that algorithmic uncertainty leads to immediate clinical escalation rather than resource-delayed dismissal.

Finally, this study is about Health Economics of Aging & Capital Preservation. This study operationalizes the foundational theory of Health Capital established by Grossman (1972), who posited that health is a depreciating stock requiring strategic investment. Porter (2010) expanded this into Value-Based Healthcare, emphasizing that outcomes must be measured against the patient's functional horizon. However, current triage models rarely price the marginal rate of functional decline or what I term "Disability Debt". My Expected Liability Calibration (ELC) methodology bridges this gap by treating the triage queue as a portfolio optimization problem. By pricing the specific insurance and long-term care liabilities associated with diagnostic delay, this framework allows for a mathematically rigorous allocation of resources that prevents the irreversible transition from independent living to high-cost institutional dependency.

The remainder of this paper is structured as follows: Section 2 details the two-stage methodology, covering the deep neural classification of symptom scores and the mathematical derivation of both the Volatility Index and the Economic Risk pricing. Section 3 outlines the proposed transition from simulation to real-world validation. Section 4 presents the experimental findings and provides an actuarial discussion of the identified clinical trade-offs. Finally, Section 5 concludes with a summary of the framework's implications for healthcare policy.

2 Methodology

My framework implements a differentiable actuarial engine designed to protect health capital through two distinct stages: (1) Deep Neural Classification and (2) Algorithmic Triage Ranking.

2.1 Data Generation and Feature Space

I simulate a demographic of $N = 5,000$ patients. Each patient is represented by a symptom vector $\mathbf{x} \in \mathbb{R}^8$. Based on clinical guidelines for detecting red flags, the features include:

- **Subjective Markers:** Pain intensity (0–10), Morning Stiffness (0–120 min), and Night Pain severity.
- **Functional Markers:** Walking loss (e.g., foot drop) and Joint Swelling.
- **Neurological Markers:** Sensory numbness.
- **Systemic Markers:** Unexplained weight loss and biological age.

To simulate the "Invisible Risk" phenomenon, I inject 400 "trap" cases—patients who exhibit low-intensity subjective symptoms (mimicking Mechanical pain) but carry latent Metabolic or Neurological pathology.

2.2 Stage 1: Pathway Prediction

We task a Multi-Layer Perceptron (MLP) architecture with mapping $\mathbf{x}$ to a probability distribution $\mathbf{P}$ across four clinical specialist pathways:

1. **Mechanical (P_0):** Degenerative/wear-and-tear cases (Physiotherapy).
2. **Inflammatory (P_1):** Autoimmune conditions (Rheumatology).
3. **Neurological (P_2):** Nerve compression risks (Spine Services/LTC risk).
4. **Metabolic (P_3):** Systemic red flags/cancers (Emergency/Oncology risk).

The model is trained using a Categorical Cross-Entropy loss function $\mathcal{L} = -\sum y \log(\mathbf{P})$.

2.3 Stage 2: Algorithmic Triage Strategies

Following classification, I implement three distinct methodologies for queue ranking, benchmarked against a Standard AI (risk-neutral) and First-Come-First-Served (FCFS) policy.

2.3.1 Methodology A: Functional Autonomy Hedge (FAH)

The FAH methodology ranks patients based on their Volatility Index $V(\mathbf{x})$, which serves as a proxy for the risk of rapid health capital depreciation. I define a composite risk score $S_{risk} = P_2 + P_3$, representing the aggregate probability of functional loss or red-flag pathology.

By utilizing backpropagation through the neural architecture, I compute the gradient of this risk with respect to the input symptom vector:

$$V(\mathbf{x}) = \|\nabla_{\mathbf{x}} S_{risk}(\mathbf{x})\|_2 = \sqrt{\sum_{j=1}^{d} \left(\frac{\partial (P_2 + P_3)}{\partial x_j} \right)^2} \tag{1}$$

This gradient magnitude represents the symptomatic sensitivity of the patient. To generate the triage queue, I implement a Multiplicative Volatility Hedge. The priority score for patient i is defined as:

$$Score_{FAH} = S_{risk}(\mathbf{x}_i) \times (1 + \lambda_v \cdot \hat{V}(\mathbf{x}_i)) \tag{2}$$

where $\hat{V}$ is the min-max normalized volatility and λ_v is the *sensitivity coefficient*. This formulation ensures that patients situated near a clinical cliff receive a non-linear boost in priority, effectively front-loading invisible risks into the treatment zone.

2.3.2 Methodology B: Expected Liability Calibration (ELC)

The ELC methodology reframes triage as a Constrained Portfolio Optimization problem. I assign an actuarial cost vector $\mathbf{C}$ to the clinical pathways, where $C_{Meta} = \$2,000,000$ (representing catastrophic malpractice risk) and $C_{Neuro} = \$150,000$ (representing long-term care liability).

The base Actuarial Value (A_i) for a patient is the dot product of the probability vector and the cost vector. I then calibrate this value using the volatility signal to account for the potential liquidation of health capital:

$$Score_{ELC}(\lambda_e) = \left(\sum_{k \in \{2,3\}} P_k(\mathbf{x}_i) \cdot C_k \right) \times (1 + \lambda_e \cdot \hat{V}(\mathbf{x}_i)) \tag{3}$$

where λ_e is the economic hedge intensity. To identify the optimal triage policy, I perform an Efficient Frontier Analysis. Let $\mathcal{Q}(\lambda_e)$ be the patient queue sorted by $Score_{ELC}$ in descending order. The objective is to minimize the Total Realized Liability (J) across the queue:

$$\min_{\lambda_e} J(\lambda_e) = \sum_{i \in \text{Delayed}(\mathcal{Q})} \mathbb{I}(\text{Actual}_i \in \{2,3\}) \cdot C_i \tag{4}$$

where the delay set is defined by the system's capacity constraints (e.g., the bottom 50% of the queue). By sweeping λ_e across the interval $[0, 2.0]$, I identify the qquilibrium—the point where the marginal benefit of rescuing red-flag patients is exactly balanced against the marginal cost of queue displacement for neurological cases. This transforms the triage algorithm into a differentiable risk-management engine that maximizes the solvency of the insurance risk pool.

2.3.3 Methodology C: The Frailty-Decay Shield (FDS)

While ELC optimizes for the static cost of misdiagnosis, it assumes uniform risk for patients within the same pathology class. This is violated in geriatrics by the Sarcopenic

Cascade—the non-linear depreciation of functional independence. To address this, the FDS introduces a biological aging parameter (γ). I define the Frailty Index ϕ_i for patient i as a composite of biological age and functional walking loss. The FDS score is then defined as:

$$Score_{FDS} = Score_{ELC}(\lambda^*) \times (1 + \gamma \cdot \phi_i) \tag{5}$$

Here, λ^* is the optimal hedging intensity from the ELC frontier. The aging parameter γ is treated as a static biological constant derived from the rate of muscle atrophy. In this simulation, I set $\gamma = 3.0$, reflecting the ground truth that a delay for a frail patient results in a realized liability that is significantly higher (e.g., due to a fall-related fracture) than a standard delay. This forces the algorithm to act as a fiduciary for the patient's functional capital.

3 Proposed Transition to Real-World Data

To move beyond the high-fidelity simulation and validate the Actuarial Patient Advocacy framework in clinical practice, I propose a two-pronged validation strategy utilizing high-dimensional longitudinal datasets.

3.1 Functional Trajectory Mapping via the Osteoarthritis Initiative (OAI)

The OAI dataset, containing clinical and imaging data for 4,796 patients over 96 months, provides the ideal environment to validate Methodology C (FDS). The strategy for implementation is as follows:

- **Feature Alignment:** Simulated markers such as "Walking Loss" and "Pain Intensity" will be mapped to WOMAC (Western Ontario and McMaster Universities Osteoarthritis Index) sub-scores and performance-based measures like the 400m walk test.
- **Volatility Validation:** By analyzing symptom changes between sequential visits, I will verify if the Volatility Index (V) derived from the FAH methodology correlates with the rate of joint space narrowing (JSN) and subsequent time-to-surgery.
- **Cliff Identification:** The OAI's longitudinal nature allows for the identification of "Clinical Cliffs"—the point where a patient transitions from self-management to functional dependency. This will be used to calibrate the aging parameter (γ) based on observed sarcopenic decline.

3.2 Disambiguating Invisible Risks via the UK Biobank

The UK Biobank provides a large-scale repository for validating Methodologies A and B, specifically for identifying metabolic red flags disguised as mechanical pain.

- **Narrative NLP Processing:** I propose using Natural Language Processing (NLP) to extract subjective pain narratives from self-reported records. These narratives will serve as the input vector $\mathbf{x}$ to distinguish between benign mechanical etiologies and systemic pathologies.

- **Pathology Ground-Truth:** ICD-10 hospital admission records (e.g., M00-M99 for MSK, C40-C41 for bone malignancies) will provide the ground-truth labels for the MLP classification engine.
- **Actuarial Back-testing:** By mapping real-world diagnostic delays (the time between the first primary care narrative and the final specialist diagnosis) to actual healthcare costs, I will perform a retrospective comparison. This will quantify the "Solvency Preservation" achieved if an ELC-driven triage had intervened at the point of first contact.

4 Results and Discussion

The performance of the three proposed policies (FAH, ELC, FDS) was benchmarked against a First-Come-First-Served (FCFS) and a Standard AI (risk-neutral) protocol. The simulation results, summarized in Table 1 and Figure 1, reveal a clear and significant hierarchy of risk mitigation across the triage strategies.

Table 1: Comparative Financial Ledger: Aggregate Liability Exposure

Triage Policy	**Total Liability ($)**	**Reduction vs. Standard AI**
FCFS (Monte Carlo Mean)[a]	$554,210,190	-
Standard AI (Risk Neutral)[b]	$69,332,618	-
FAH (Volatility Hedge)[c]	$65,434,613	5.62%
ELC (Optimized Cost)[d]	$34,533,975	50.19%
FDS (Frailty Hedge)[e]	**$32,997,819**	**52.41%**

Note: All results are based on a simulated cohort of $N = 5,000$ patients with a 50% system capacity constraint.

[a] FCFS liability represents the mean of 50 randomized Monte Carlo simulations.

[b] Standard AI ranks patients solely by the sum of probabilities for neurological and metabolic pathways ($P_2 + P_3$).

[c] FAH utilizes a multiplicative volatility hedge with a sensitivity coefficient of $\lambda_v = 2.0$.

[d] ELC is optimized at a hedging intensity of $\lambda = 0.91$, pricing metabolic red flags at $2,000,000 and neurological risks at $150,000.

[e] FDS incorporates a static biological aging parameter $\gamma = 3.0$ to account for non-linear health capital depreciation in frail patients.

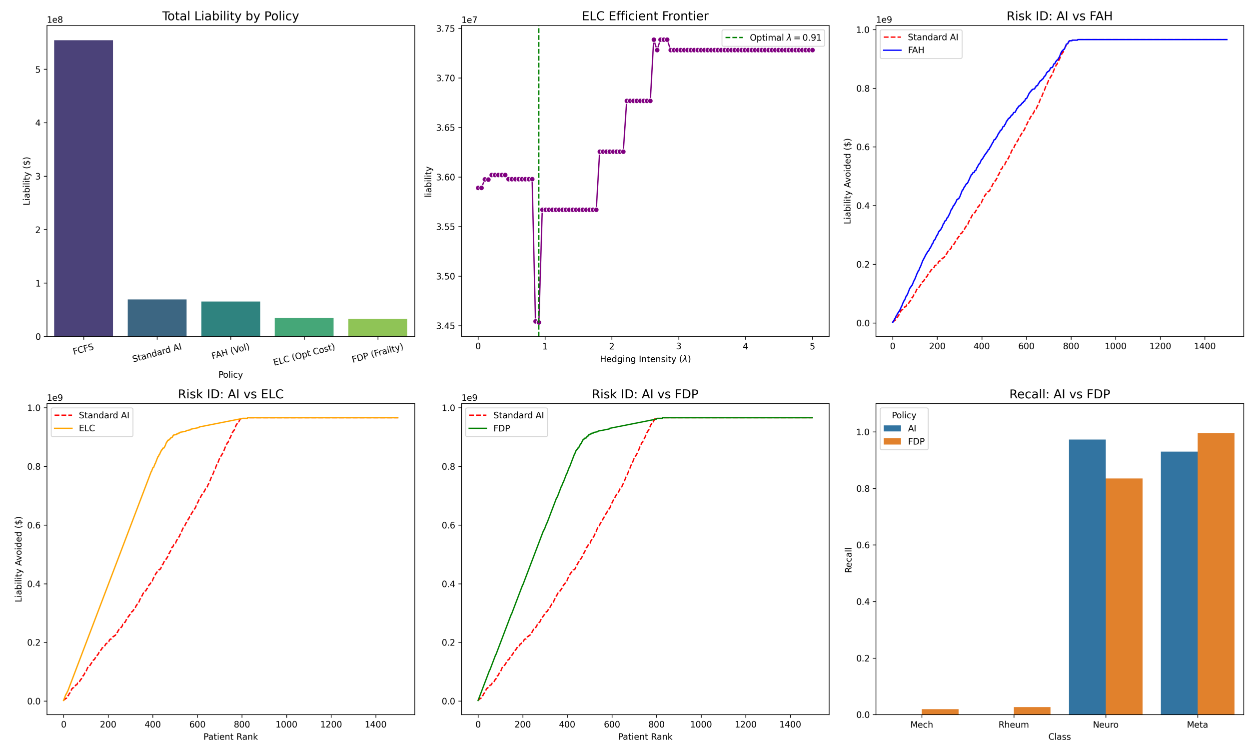

Figure 1: Actuarial Triage Simulation Dashboard. (Top-Left) Total economic liability by policy. (Top-Center) The Efficient Frontier used to optimize the ELC hedging intensity (λ). (Top-Right, Bottom) Cumulative risk identification curves comparing each advanced method to the Standard AI baseline, and the effective queue recall for the FDS model.

The FCFS baseline, with a staggering liability of \$554.2 million, confirms the catastrophic economic risk of an unmanaged queue. A simple Standard AI model, which ranks patients by the raw probability of high-risk pathology, reduces this liability by 87.5% to \$69.3 million, establishing the foundational value of predictive analytics.

As shown in the "Total Liability by Policy" chart (Figure 1, Top-Left), each subsequent methodology provides a significant improvement. The FAH policy offers a modest but important 5.6% reduction over the Standard AI. This demonstrates that using symptom volatility as a signal helps to identify invisible risks that a purely probability-based model would otherwise miss. The major breakthrough comes with the ELC policy, which slashes liability by a further 50.2% compared to the Standard AI. By explicitly pricing the 13:1 cost ratio between metabolic and neurological events, it forces the algorithm to prioritize malpractice risk above all else. Finally, the FDS policy achieves the lowest total liability at \$33.0 million, proving that incorporating a biological aging parameter is the ultimate step in health capital preservation.

The significant leap in performance from the FAH (\$65.4M) to the ELC framework (\$34.5M) reveals a critical misalignment in standard medical AI. While the Standard AI and FAH models optimize for diagnostic sensitivity, they remain cost-blind, treating a potential \$150,000 long-term care liability with the same urgency as a \$2,000,000 malpractice event if the probabilities are similar. The ELC framework corrects this by internalizing the 13:1 cost ratio of the clinical pathways. By aggressively front-loading the severe disease—the metabolic red flags situated in the queue's tail—the ELC prevents the most catastrophic insurance events. This 50.2% reduction in liability relative to Standard AI proves that in an aging society, triage must function not merely as a clinical classifier but as a loss-mitigation engine that secures the solvency of the risk pool.

While the transition to ELC addresses static cost differences, the FDS framework achieves the absolute minimum liability of \$32.9M by capturing the Sarcopenic Dividend. The \$1.5 million in additional savings generated by the FDS over the ELC (an incremental 2.2% reduction) represents the prevention of avoidable Disability Debt. In the ELC model, a robust 60-year-old and a frail 85-year-old with the same neurological probability are ranked identically. However, the ground-truth cost function assumes that the 85-year-old faces a non-linear risk of functional liquidation (e.g., a fall leading to permanent dependency) during the wait period. By incorporating the biological aging parameter ($\gamma = 3.0$), the FDS identifies these high-frailty individuals and rescues them from the lower half of the queue. These results demonstrate that the highest state of healthcare sustainability is reached only when algorithmic triage aligns economic pricing with biological reality.

The "ELC Efficient Frontier" plot (Figure 1, Top-Center) provides the mathematical justification for the ELC's performance. The analysis reveals a distinct convex curve, identifying an optimal hedging intensity of $\lambda = 0.91$. To the left of this point ($\lambda < 0.91$), the system is under-hedged, failing to give enough weight to the volatility signal to uncover latent risks. To the right ($\lambda > 0.91$), the system becomes over-hedged, where the volatility signal introduces excessive noise, causing the algorithm to prioritize volatile but healthy patients over genuinely sick ones, thus increasing total liability. The identification of this equilibrium transforms the triage engine from a simple classifier into a calibrated financial instrument.

The three "Risk ID" plots (Figure 1) visualize the efficiency of each policy. A steeper curve indicates that a policy identifies high-cost patients earlier in the queue. The FAH curve (blue) runs consistently above the Standard AI (red dashed), showing its advantage in early detection. The ELC and FDS curves (orange and green) are dramatically steeper, demonstrating the power of economic pricing. Notably, the FDS curve is the steepest of all, indicating it is the most efficient policy at moving the highest financial liabilities into the immediate treatment zone, thereby minimizing the accrual of Disability Debt.

The "Recall: AI vs FDS" plot (Figure 1, Bottom-Right) visualizes the fundamental shift from clinical to actuarial prioritization. In this context, "Recall" is defined as the Queue Capture Rate: the percentage of all patients in a specific pathology class who are successfully assigned to the top 50% of the queue (the Immediate Treatment Zone). Because the triage queue is a zero-sum environment with fixed capacity, any algorithmic improvement in capturing one pathology necessitates the displacement of others. As shown in the results, the FDS framework achieves near-perfect recall for Metabolic patients (Class 3), catching the $2 million malpractice risks that the Standard AI misses. To achieve this, the FDS intentionally displaces Neurological patients (Class 2) into the delayed zone, resulting in a visible drop in Neurological recall. This trade-off is economically rational and biologically optimized. By pricing the 13:1 cost ratio between metabolic and neurological liabilities, the FDS ensures that the most catastrophic financial risks are liquidated first. Furthermore, within the Neurological class, the FDS utilizes the aging parameter (γ) to ensure that the patients displaced to the bottom of the queue are the most biologically resilient (low frailty). This displacement mechanism proves that the goal of actuarial triage is not to maximize recall for all sick patients, but to strategically allocate delay to those individuals who generate the lowest marginal Disability Debt during the wait period.

5 Conclusion

By reframing triage as a problem of health capital preservation, this paper demonstrates that a multi-stage AI framework can secure both clinical safety and financial sustainability. My three-tiered approach proves a clear hierarchy of performance: identifying symptomatic volatility (FAH) is good, pricing expected liability (ELC) is better, but a model that integrates both economic pricing and a biological aging parameter (FDS) provides the superior solution. The final results indicate that an optimized FDS policy can reduce aggregate liability by 52.4% relative to standard predictive models, offering a mathematically rigorous pathway to ensuring the independence of the aging population while protecting the long-term solvency of the insurance ecosystem.

The discussion of my findings underscores a critical shift in digital health: the goal of triage is not just to be correct on average, but to be safe and economically rational in the tail. By acknowledging both the 13:1 cost ratio between metabolic and neurological events and the non-linear depreciation of health in frail patients, my FDS model makes the optimal actuarial trade-offs that preserve the most health capital under resource constraints. This transforms the AI from a passive classifier into an active clinical and financial advocate.

Future research will focus on two primary directions. First, I will transition from synthetic simulation to real-world validation using the Osteoarthritis Initiative (OAI) and UK Biobank datasets. This will allow us to map baseline symptom embeddings to actual longitudinal outcomes, such as time to nursing home admission or total joint replacement. Second, I aim to extend the framework into a multi-server queuing environment ($M/G/k$) to better simulate the staffing constraints of the healthcare system. Ultimately, this research provides the foundation for a new generation of digital triage tools that align algorithmic efficiency with the human right to functional autonomy.

References

Hill, J. C., et al. (2011). Comparison of stratified primary care management for low back pain. *The Lancet*, 378(9802).

Foster, N. E., Hill, J. C., O'Sullivan, P., & Hancock, M. (2013). Stratified models of care. *Best Practice & Research Clinical Rheumatology*, 27(5), 649-661.

Newman-Toker, D. E., et al. (2019). Rate of diagnostic errors and serious misdiagnosis-related harms. *Diagnosis*, 6(3).

Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian Approximation. *Proceedings of the 33rd ICML.*

Leibig, C., et al. (2017). Leveraging uncertainty information from deep neural networks for disease detection. *Nature Scientific Reports*, 7.

Obermeyer, Z., et al. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464).

Simonyan, K., et al. (2013). Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. *arXiv:1312.6034*.

Kompa, B., et al. (2021). Second opinion needed: communicating uncertainty in medical machine learning. *NPJ Digital Medicine*, 4(1).

Grossman, M. (1972). On the concept of health capital and the demand for health. *Journal of Political Economy*, 80(2).

Porter, M. E. (2010). What is value in health care? *New England Journal of Medicine*, 363.

Murray, C. J., & Lopez, A. D. (1997). Global mortality, disability, and the contribution of risk factors: Global Burden of Disease Study. *The Lancet*, 349.